arXiv:1507.00832v1 [math.ST] 3 Jul 2015

The Efficiency of Density Deconvolution

Stefan Wager*

Department of Statistics
Stanford University
swager@stanford.edu

Abstract

The density deconvolution problem involves recovering a target density g from a sample that has been corrupted by noise. From the perspective of Le Cam's local asymptotic normality theory, we show that *non-parametric* density deconvolution with Gaussian noise behaves similarly to a low-dimensional *parametric* problem that can easily be solved by maximum likelihood. This framework allows us to give a simple account of the statistical efficiency of density deconvolution and to concisely describe the effect of Gaussian noise on our ability to estimate g, all while relying on classical maximum likelihood theory instead of the kernel estimators typically used to study density deconvolution.

KEYWORDS. Adaptive minimaxity, Gaussian sequence model, local asymptotic normality, relative efficiency.

1 Introduction

Suppose that we observe n samples X_i drawn from a hierarchical model

$$\mu_i \overset{\text{iid}}{\sim} g(\cdot), \quad X_i = \mu_i + \varepsilon_i, \quad \text{with} \quad \varepsilon_i \overset{\text{iid}}{\sim} \mathcal{N}(0, 1), \tag{1}$$

and our goal is to estimate the unknown density $g(\cdot)$. This problem, sometimes called the density deconvolution problem, is remarkably hard in terms of asymptotic statistical criteria. For example, as shown by Carroll and Hall [1988] and Fan [1991], if we assume a non-parametric setup where g is only known to have a Lipschitz-continuous k-th derivative, then the minimax error rate for estimating g under the integrated squared error loss decays as $\log(n)^{-(k+1)}$.

The density deconvolution problem is traditionally studied using kernel methods. The motivation for kernel estimators is that they often achieve adaptive

*I am deeply grateful to Brad Efron for many enlightening conversations as well as his continual encouragement, and to Dave Donoho and Will Fithian for several helpful comments and suggestions. This work was supported by a B. C. and E. J. Eaves Stanford Graduate Fellowship.

minimaxity over natural regularity classes for g; for example, they in fact achieve the optimal rates of Carroll and Hall [1988] and Fan [1991] described above. The properties of kernel density deconvolution have been analyzed by several authors, including Butucea and Comte [2009], Carroll and Hall [2004], Comte and Lacour [2011], Efromovich [1997], Fan and Koo [2002], Hall and Lahiri [2008], Hall and Meister [2007], Stefanski and Carroll [1990], Wand [1998], and Zhang [1990].

Despite the prolific literature devoted to them, however, kernel methods do not necessarily yield a fully satisfying theory of the statistics of density deconvolution. In particular, Efron [2014a,b] proposed a simple maximum likelihood approach to density deconvolution that performs qualitatively better than kernel methods on several realistic scientific tasks. Thus, it appears that while kernel methods are nearly optimal in terms of the standard asymptotic criteria used in the literature, they are not always optimal in practical applications. This suggests a need for a new optimality theory for density deconvolution.

The goal of this paper is to move us towards such a theory. We begin with a close analysis of Efron's method, which involves estimating the unknown density g by maximum likelihood with a p-parameter model

$$g_\eta(\mu) = g_0(\mu) \exp\left[\eta \cdot T(\mu) - \psi(\eta)\right], \tag{2}$$

where $T(\mu)$ is some carefully chosen p-parameter statistic and $\psi(\cdot)$ is the log-partition function. Given that Efron's method involves *parametric* maximum likelihood estimation, it may appear surprising that this method would work well in a *non-parametric* setup where we only know that the target density g belongs to some regularity class. However, we show that maximum likelihood estimation in the model (2) with an appropriate choice of T has adaptive minimaxity properties that are reminiscent of those enjoyed by kernel estimators. Moreover, because Efron's method relies on maximum likelihood instead of the ad-hoc kernel inversion procedure used by classical methods, we may also hope for the method to be well-behaved in a wide variety of practical applications.

We build our analysis around a *relative efficiency* criterion, which measures the information loss for estimation in the model (2) we incur from the noise ε in (1). Our main result is that, given a specified carrier g_0, there exists a low-dimensional statistic T that captures essentially all the information contained in the X-sample for estimating local perturbations to g_0. Any model of the form (2) that tries to use a higher-dimensional parametrization will have some bad contrasts that could have been accurately estimated from clean observations μ_i, but become effectively impossible to estimate from the noised observations X_i. If g_0 is Gaussian, the number of samples needed to accurately estimate a p-parameter family of the form (2) scales exponentially in p, and this bound holds uniformly over any possible choice of the statistic T.

At face value, our results can be interpreted as a companion to the classical hardness results of Carroll and Hall [1988] and Fan [1991]: our analysis implies that it is impossible to accurately estimate any parametric family of the form (2) in terms of relative efficiency if p is even moderately large. Thus, no matter

how clever we may be, we cannot use Efron's density deconvolution model to efficiently learn a rich model for g.

However, from a decision-theoretic point of view, our result can also be interpreted in a more optimistic light. Because high-dimensional models of the form (2) are effectively impossible to estimate, we lose almost nothing by just using a low dimensional model. Thus, the non-parametric density deconvolution problem effectively reduces to parametric inference for Efron's model (2) with an appropriate low-dimensional parametrization T.

Our results take on a particularly simple form when the carrier g_0 is Gaussian. In this case, the "optimal" statistics are polynomials in μ, yielding the class of estimators

$$\hat{g}_\eta^p(\mu) = \exp\left[\sum_{j=1}^{p} \eta_j \mu^j - \psi_p(\eta)\right], \tag{3}$$

where p is a tuning parameter. As we will show, given an appropriate (usually small) choice of p, parametric inference for this deceptively simple model is nearly equivalent to optimal non-parametric inference for g.[1] More specifically, this class of estimators is asymptotically nearly adaptively minimax under Kullback-Leibler (KL) loss for estimating local perturbations of g_0

$$g(\mu) = g_0(\mu)\exp\left[\tau(\mu) - \psi_\tau\right], \quad g_0(\mu) = \frac{1}{\sigma}\varphi\left(\frac{\mu}{\sigma}\right), \tag{4}$$

where $\tau(\cdot)$ is a tilting function satisfying regularity conditions detailed in Section 3. Moreover, the efficiency shortfall of the best estimator of the form (3) relative to the minimax estimator can be bounded by a small explicit constant that is, for typical parameter values, on the order of 2.

In summary, this paper introduces a relative efficiency criterion and a local perturbation model that, together, enable us to shed new light on the classical problem of density deconvolution. We show that the exponential family method (2) cannot get around standard non-parametric impossibility results and, in particular, does not allow us to estimate rich models for $g(\cdot)$. However, the method does an excellent job at extracting all the available information from the X-sample, all while allowing for straight-forward estimation and inference. Thus, despite its simple form, Efron's parametric density deconvolution method yields simple estimators whose excellent practical performance is firmly grounded in asymptotic minimaxity theory.

[1] In all our experiments, we use $p = 4$. In practice, p could also be selected by cross-validation.

1.1 Theoretical Setup: Local Deconvolution

Throughout this paper, we assume that we observe n samples X_i drawn from a hierarchical model

$$\mu_i \overset{\text{iid}}{\sim} g(\cdot), \; X_i = \mu_i + \varepsilon_i, \; \varepsilon_i \overset{\text{iid}}{\sim} \mathcal{N}(0, 1); \text{ thus,} \tag{5}$$

$$X_i \overset{\text{iid}}{\sim} f(\cdot), \; f(x) = (\varphi * g)(x), \tag{6}$$

where f denotes the marginal density of the observations X generated according to the model (1) and φ is the standard Gaussian density. The setup induced by (5) and (6) is classical; however, the analysis techniques we use to understand this model are rather different from the ones used by, e.g., Carroll and Hall [1988], Efromovich [1997], or Fan [1991].

Unlike other authors who model $g(\cdot)$ as a fixed member of a regularity class $\mathcal{R}$ that can be estimated with increasing accuracy as the sample size n grows to infinity, we study a sequence of ever-shrinking perturbations of a known carrier density g_0:

$$g^{(n)}(\mu) = g_0(\mu) \exp\left[\frac{1}{\sqrt{n}}\tau(\mu) - \psi_n\right], \tag{7}$$

where $\tau(\mu)$ belongs to an appropriate regularity ellipsoid. Notice that, as n gets large, estimating τ does not necessarily get easier because the deviance between g_0 and $g^{(n)}$ decays as $1/\sqrt{n}$. As we will show formally in Section 3, the problem of estimating $g^{(n)}$ in (7) under a re-scaled deviance loss

$$L_n\left(\hat{g}^{(n)}\right) = n\, D_{KL}\left(g^{(n)}, \hat{g}^{(n)}\right) = n \int g^{(n)}(\mu) \log\left(\frac{g^{(n)}(\mu)}{\hat{g}^{(n)}(\mu)}\right) d\mu. \tag{8}$$

converges to a *locally asymptotically normal* experiment in the sense of Le Cam [1960]. The formalism provided by the sequence of problems (7) thus enables us to study density deconvolution from the perspective of classical maximum likelihood asymptotics.

Although the estimation problem in the local perturbation model (7) may look very different from the classical estimation problem with $g \in \mathcal{R}$ for some fixed regularity set $\mathcal{R}$, it appears that studying the former can help us learn about the latter. First of all, we find a universality phenomenon, where maximum likelihood estimation in the model (3) is simultaneously nearly minimax for any Gaussian carrier $g_0 = \varphi_\sigma$, and so we can carry out practical data analysis with a simple default model for g. Second, our analysis recovers familiar qualitative aspects of the theory of Carroll and Hall [1988] and Fan [1991], such as exponential blow-ups in the number of samples required to estimate more complex models; see, e.g., Theorem 5. Thus, insofar as maximum-likelihood estimation in our local perturbation model is amenable to exact asymptotic analysis, it appears that the model (7) is a useful theoretical tool for gaining new insights about the density deconvolution problem.

We begin our analysis by developing a theory of relative efficiency for density deconvolution in Section 2, with the goal of establishing lower bounds for the

error rate of any p-parameter model of the form (2). In Section 3, we complement this worst-case picture by establishing near-minimax properties of the estimator (3) for estimating local perturbations of Gaussian densities. Our proof technique is built on the local asymptotic normality theory of Le Cam [1960] combined with Pinsker's theorem for the Gaussian sequence model. Finally, in Section 4, we discuss practical differences between kernel density deconvolution and Efron's method, and the respective optimality theories that justify each method. All proofs are provided in the appendix.

Our minimaxity proof, which first reduces a continuous problem to a Gaussian sequence model and then applies Pinsker's theorem, fits into a rich literature on solving non-parametric problems via Gaussian estimation; see, e.g., Brown et al. [2004], Brown and Low [1996], Efromovich and Samarov [1996], Golubev et al. [2010], Johnstone [2011], and Nussbaum [1996]. The problem of density estimation using a log-spline model for g of the form (2) has also been considered by Koo [1999] and Koo and Chung [1998]; however, their approach more closely follows that traditional ideas of Fan [1991] and others. Finally, we note that the non-parametric maximum likelihood problem of estimating $f = \varphi * g$ has been studied by, among others, Jiang and Zhang [2009], Koenker and Mizera [2014], and Zhang [2009]. Efron's method also induces a natural estimator $f_{\hat{\eta}} = \varphi * g_{\hat{\eta}}$, but we do not study its properties here.

Remark: Cyclic Convolution To avoid technical difficulties relating to integrability over unbounded domains, we follow convention and study a *cyclic* convolution model [e.g., Efromovich, 1997]. More specifically, we assume that our observations are within a bounded interval $\Omega_M = [-M, M]$, and that we have a cyclic convolution operator

$$K_M : \Omega_M \times \Omega_M \to \mathbb{R}_+, \quad K_M(\mu, x) = \sum_{j=-\infty}^{\infty} \varphi(x - \mu + 2jM). \tag{9}$$

We also frequently use the cyclically wrapped-around Gaussian density with variance σ^2, i.e., $\varphi_\sigma^M(x) = \sum_{j=-\infty}^{\infty} \sigma^{-1} \varphi((x + 2jM)/\sigma)$ with $x \in \Omega_M$, as the carrier g_0. In our results, M should be thought of as large enough that $K_M(\mu, x) \approx \varphi(x - \mu)$ over the range of the data; formally, we will seek to state results in a limit with $M \to \infty$.

2 Relative Efficiency and Most Stable Families

The observations X_i generated by the model (1) are noisy measurements of clean draws μ_i from g. What is the information loss due to this extra noise? Or, in other words, if we can estimate η to a given accuracy using n_μ samples μ drawn directly from g, how many samples n_X would we need to estimate η to the same accuracy using only X samples? Although this question may not at first appear directly related to our topic of interest, answering it will prove to be

helpful in guiding us towards good choices for the statistic T in Efron's model (2) and in proving lower bounds for the error rate of any parametric model.

In Section 2.1 below, we define a *relative efficiency coefficient* that will let us make the above question precise. Using this formalism, we then derive the "best" families of the form (2) in terms of relative efficiency in Section 2.2. Finally, in Sections 2.3 and 2.4 we interpret these relative efficiency results for both Gaussian and non-Gaussian carriers g_0 respectively.

2.1 The Relative Efficiency Coefficient

If we had access to samples μ_i drawn directly from a density g_η of the form (2), then under standard regularity conditions maximum likelihood estimation would yield an asymptotically normal estimator $\hat{\eta}_\mu$ with

$$\sqrt{n}\left(\hat{\eta}_\mu - \eta\right) \to \mathcal{N}\left(0, \mathcal{I}_\mu^{-1}\left(\eta\right)\right), \quad \mathcal{I}_\mu\left(\eta\right) = -\mathbb{E}_\eta\left[\frac{\partial^2}{\partial \eta^2} \log g_\eta\left(\mu\right)\right], \tag{10}$$

where $\mathcal{I}_\mu\left(\eta\right)$ is the Fisher information for estimating η carried by the μ-samples. In our setup, however, we only have access to samples X drawn from (6); the Fisher information for the maximum-likelihood estimator $\hat{\eta}_X$ is then reduced to

$$\mathcal{I}_X\left(\eta\right) = -\mathbb{E}_\eta\left[\frac{\partial^2}{\partial \eta^2} \log f_\eta\left(X\right)\right]. \tag{11}$$

Given this setup, we define the efficiency of $\hat{\eta}_X$ relative to $\hat{\eta}_\mu$ as

$$\rho_\eta\left(T\right) = \inf_{a \in \mathbb{R}^p,\, a \neq 0}\left\{\frac{a^\top \mathcal{I}_X\left(\eta\right) a}{a^\top \mathcal{I}_\mu\left(\eta\right) a}\right\}, \tag{12}$$

and argue that the relative efficiency coefficient ρ provides a natural measure of information loss due to the noise ε in (1).

If we wanted to measure relative efficiency in a univariate family (2) with tilting function $t : \mathbb{R} \to \mathbb{R}$, then the natural measure of relative efficiency is

$$\rho_\eta\left(t\right) = \lim_{n \to \infty} \frac{\mathrm{Var}\left[\hat{\eta}_\mu\right]}{\mathrm{Var}\left[\hat{\eta}_X\right]} = \frac{\mathcal{I}_X\left(\eta\right)}{\mathcal{I}_\mu\left(\eta\right)}, \tag{13}$$

which measures exactly the increase in sample size required for accurate estimation using the X-sample instead of the μ-sample. As we can easily verify, the multivariate relative information coefficient (12) is simply the worst-case relative efficiency for any univariate subfamily of (2):

$$\rho_\eta\left(T\right) = \inf\left\{\rho_\eta\left(t\right) : t\left(\mu\right) = a \cdot T\left(\mu\right),\, a \in \mathbb{R}^p\right\}. \tag{14}$$

This connection provides a first motivation for the definition (12). The coefficient ρ_η can also be viewed as a natural extension to the classical E-optimal criterion for experimental design [Ehrenfeld, 1955]: if T is scaled such that $\mathcal{I}_\mu = I_{p \times p}$, then $\rho_\eta\left(T\right)$ is the minimum eigenvalue of $\mathcal{I}_X$—which is exactly the criterion that the E-optimal designs maximize.

The following result derives a simple functional form for the relative efficiency coefficient. We note a striking similarity between the formula (15) and the results of Louis [1982] for maximum likelihood estimation with missing data. This similarity is not an accident: we could also interpret the density deconvolution problem as a one where we would would have wanted to observe $\mu = X - \varepsilon$, but ε is missing.

We also establish a key theoretical property of ρ, namely that it is transformation invariant: if we apply any invertible linear transformation Q to T, the value of ρ remains unchanged. This transformation invariance provides further evidence that ρ is a "natural" measure for understanding the difficulty of density deconvolution. We note in particular that the Fisher information $\mathcal{I}_X$ is not transformation invariant.

Lemma 1. *The relative efficiency coefficient (12) is transformation invariant, in the sense that $\rho_0(T) = \rho_0(QT)$ for any invertible linear transformation Q. For a univariate statistic, the relative efficiency coefficient can be written as*

$$\rho_\eta(t) = \frac{\mathrm{Var}_\eta\left[\mathbb{E}\left[t(\mu) \mid X\right]\right]}{\mathrm{Var}_\eta\left[t(\mu)\right]}. \tag{15}$$

2.2 Deriving the Most Stable p-dimensional Family

Given our notion of relative efficiency defined above, it is natural to ask whether there exist optimal p-dimensional statistics T in terms of this criterion. Perhaps surprisingly, we will show that not only do such optimal statistics exist, but they are in general quite easy to compute and can give us guidance for practical data analysis. Formally, we define the optimal statistics as solutions to the optimization problem

$$\Gamma_p(g_0) = \mathrm{argmax}_{T:\Omega\to\mathbb{R}^p}\left\{\rho_0(T)\right\}, \tag{16}$$

where $g_\eta(\mu) = g_0(\mu)\exp\left[\eta\cdot T(\mu) - \psi(\mu)\right]$ is defined in terms of some known carrier. Notice that our definition of Γ in terms of relative efficiency at $\eta = 0$ is without loss of generality, since we could always just use $g_\eta(\mu)$ as our carrier if this condition did not hold.

In order to solve the problem (16), we begin with a technical result for computing the multivariate relative efficiency coefficient. In this section, we will assume that the X and μ are distributed over a compact interval Ω, and that X is noised by a generic convolution operator $K(\mu, x)$.

Lemma 2. *Let T be a p-dimensional statistic; let $g(\cdot) = g_\eta(\cdot)$ be defined as in (2) and let f be the marginal density of the observations X. Then*

$$\rho_\eta(T) = \rho_\eta(a^*\cdot T), \ \textit{with} \tag{17}$$

$$a^* = \underset{\|a\|_2=1}{\mathrm{argmin}}\left\{\frac{a^\top\int_\Omega T(\mu)\,T^\top(\mu)\,K^2(x,\,\mu)\,g^2(\mu)\,f^{-1}(x)\,dx\,d\mu\;a}{a^\top\int_\Omega T(\mu)\,T^\top(\mu)\,g(\mu)\,d\mu\;a}\right\}. \tag{18}$$

We note that, in practice, the optimization problem described in (18) below can be efficiently solved by taking $a^* = Q_T^{-1} b^* / ||Q_T^{-1} b^*||_2$, where b^* is the eigenvector corresponding to the smallest eigenvalue of $Q_T^\top M_T Q_T$, and the matrices M_T and Q_T are defined by

$$M_T = \int_\Omega T(\mu) T^\top(\mu) K^2(x, \mu) g^2(\mu) f^{-1}(x) \, dx \, d\mu, \tag{19}$$

$$Q_T^\top \int_\Omega T(\mu) T^\top(\mu) g(\mu) \, d\mu \; Q_T = I_{p \times p}. \tag{20}$$

With this result in hand, we are now ready to find the most favorable statistic $\Gamma_p(g_0)$. Our construction hinges around eigenfunctions of the linear operator

$$P_g(\mu_1, \mu_2) = \int_\Omega \sqrt{g(\mu_1)} K(\mu_1, x) f^{-1}(x) K(x, \mu_2) \sqrt{g(\mu_2)} \, dx, \tag{21}$$

where f is as usual defined as $f = \varphi * g$. We can verify that the top eigenfunction of P_g is given by $\sqrt{g(\cdot)}$ and has eigenvalue 1; the p subsequent eigenvectors then generate the most favorable p-dimensional family for density deconvolution.

Theorem 3. *Suppose that both the carrier density $g_0 : \Omega \to \mathbb{R}_+$ and the kernel $K : \Omega^2 \to \mathbb{R}_+$ are continuous and bounded away from 0, and that P_{g_0} as defined in (21) is a compact operator over the space $L_2(\Omega)$ of square-integrable functions over Ω. Then, P_{g_0} admits a spectral decomposition $\zeta_1, \zeta_2, \ldots$, where the first eigenfunction $\zeta_1(\mu) = \sqrt{g_0(\mu)}$ has an eigenvalue 1. Moreover, the p-dimensional exponential family of the form (2) with statistics*

$$T_j(\mu) = \frac{1}{\sqrt{g_0(\mu)}} \zeta_{j+1}(\mu) \tag{22}$$

is a most favorable p-dimensional family in the sense of (16), and the relative efficiency coefficient corresponds to the $p+1$-st eigenvalue of P_{g_0}. Finally, if the spectrum of P_{g_0} does not have repeated eigenvalues, the most favorable family $\Gamma_p(g_0)$ is unique up to scaling and rotation.

In other words, Theorem 3 establishes a direct link between the difficulty of learning rich perturbation models around g_0, and the decay rate of the spectrum of the linear operator P_{g_0}: the best-case efficiency for learning a p-parameter model depends on the p-th non-trivial eigenvalue of P_{g_0}.

As a corollary to this result, we also see that if the relative information for learning the most favorable p-parameter family is small, then the model (2) with statistics $\Gamma_p(g_0)$ provides a good X-space approximation to any local perturbation of g_0. Recall that $D_{KL}(f, f')$ effectively measures the power of the X-sample likelihood-ratio test for distinguishing f from f'; thus, the result below implies that if $\lambda_{p+2}(P_{g_0})$ is small, it is statistically impossible to detect any deviations from the most favorable p-parameter family using only X-samples.

Corollary 4. *Suppose we have a data-generating function of the form*

$$\mu_1, \ldots, \mu_n \sim g_\tau^{(n)}(\mu) = g_0(\mu) \exp\left[\frac{\tau(\mu)}{\sqrt{n}} - \psi\left(\frac{\tau}{\sqrt{n}}\right)\right] \tag{23}$$

for some square-integrable function τ satisfying $\int_{\mathbb{R}} \tau^2(\mu)\, g_0(\mu)\, d\mu \le C^2$, and write $f_\tau^{(n)} = \phi * g_\tau^{(n)}$. Then, under the conditions of Theorem 3, there exists another tilting function $\tau^{(p,*)}$ in the span of the most-favorable p-dimensional family $\Gamma_p(g_0)$ that can closely approximate $f_\tau^{(n)}$:

$$\tau^{(p,*)} = \sum_{j=1}^{p} \gamma_j^* T_j(\mu), \quad \lim_{n\to\infty} n\, D_{KL}\left(f_\tau^{(n)}, f_{\tau^{(p,*)}}^{(n)}\right) \le \frac{1}{2} C^2 \lambda_{p+2}(P_{g_0}), \tag{24}$$

where $\lambda_{p+2}(P_{g_0})$ denotes the $p+2$-nd eigenvalue of the linear operator defined in (21).

2.3 The Hardness of Local Deconvolution near Gaussian Carriers

The relative efficiency bound given in Theorem 3 is quite general, but understanding its implications for practical data analysis may not be trivial at first glance. If we are willing to take the carrier density g_0 to be Gaussian with variance σ^2, however, the quantities defined in the statement of Theorem 3 take on simple and interpretable forms. In the result below, we assume that we are doing cyclic deconvolution over the domain $\Omega_M = [-M, M]$ for some large M; this lets us state a clean result while avoiding integrability concerns over unbounded domains.

Theorem 5. *Suppose that the carrier $g_0(\mu) = \varphi_\sigma^M(\mu)$ is the Gaussian with variance σ^2 wrapped over the interval $\Omega_M = [-M, M]$, and that we are in the cyclic convolution setup over Ω_M described in Section 1.1. Then, for any p-parameter statistic T,*

$$\rho_0(T) \le \frac{1}{(1+\sigma^{-2})^p} + o_M(1), \tag{25}$$

where the residual term $o_M(1)$ becomes negligible as M gets large. Moreover, this bound is satisfied by the model (2) whose statistics are the normalized Hermite polynomials

$$\Gamma_j(\mu) = \frac{1}{\sigma} H_j\left(\frac{\mu}{\sigma}\right), \quad H_j(\mu) = \frac{1}{\sqrt{j!}} \exp\left(\frac{\mu^2}{2}\right)\left(\frac{\partial}{\partial\mu}\right)^j \exp\left(-\frac{\mu^2}{2}\right), \tag{26}$$

for which $\mathcal{I}_\mu = I_{p\times p}$, and $\mathcal{I}_X$ is diagonal with entries $\left(1+\sigma^{-2}\right)^{-1}, \ldots, \left(1+\sigma^{-2}\right)^{-p}$.

In other words, Theorem 5 implies that, if we are trying to solve the problem (2) with a p-dimensional statistic T scaled to have $\mathcal{I}_\mu(0) = I_{p\times p}$, then we are necessarily faced with a difficult 1-dimensional sub-family with information

bounded by $(1+\sigma^{-2})^{-p}$. In particular, it is impossible to accurately distinguish local alternatives of 0 for η with less than $n \sim (1+\sigma^{-2})^p$ observations, i.e., regardless of our choice of T, the number of samples required for accurate estimation scales exponentially in p.

Qualitatively, this exponential scaling in p provides a direct analogue to the logarithmic convergence rates for kernel density estimation derived by Carroll and Hall [1988] and Fan [1991]: If we work in an asymptotic regime where we incur a bias on the order of $p^{-\zeta}$ from learning a model with only p degrees of freedom, then our result suggests that—heuristically—the error rate cannot decay faster than $\log_{1+\sigma^{-2}}(n)^{-\zeta}$. The reason this analogy is only heuristic is that our theory uses a local perturbation model, whereas that of Carroll and Hall [1988] and Fan [1991] is global. The connection between the two theories is however encouraging, in that it suggests that our local perturbation model enables us to get a nuanced grasp of the statistics on density deconvolution without changing the fundamental nature of the problem.

An important consequence of Theorem 5 is that, for all $p \geq 2$ and any $\sigma > 0$, the family of distributions attaining the bound (25), namely

$$g_\eta(\mu) = \frac{1}{\sigma}\varphi\left(\frac{\mu}{\sigma}\right)\exp\left[\sum_{j=1}^{p}\eta_j H_j\left(\frac{\mu}{\sigma}\right) - \psi(\eta)\right], \tag{27}$$

is equivalent to the family (3) after re-parametrizing η. In other words, we find that the polynomial log-density model (3) is universally the most-favorable family for density deconvolution near Gaussian carriers. Moreover, as a direct consequence of Corollary 4, we see that this family can be used to closely approximate any local perturbation to a Gaussian.

In Section 3, we will provide further justification for the estimator (27) by establishing minimax-optimality properties. Before doing so, however, we pause to study the implications of Theorem 3 in the case when g_0 is not Gaussian.

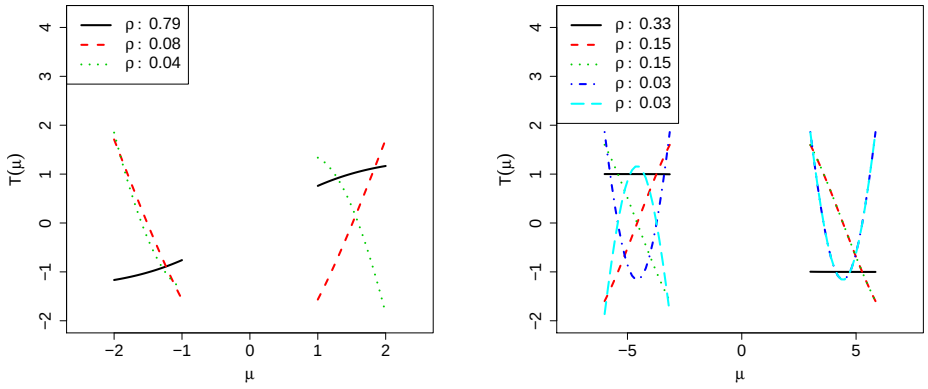
2.4 Local Density Deconvolution with Non-Gaussian Carriers

We can also use Theorem 3 to compute most favorable families for estimating perturbations to non-Gaussian carriers g_0. In this case, we can no longer derive closed-form expressions for them; however, we can still proceed numerically.

We begin by examining a "two towers" model, also considered by Efron [2014b]. We vary the scale of the carrier

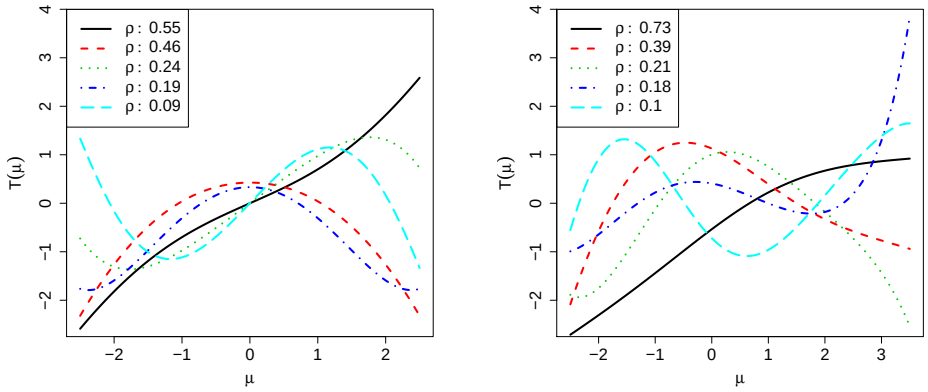
$$g_0^{(TT,\,1)}(\mu) = \frac{1}{2}\,1\left(\{1 \leq |\mu| \leq 2\}\right) \quad \text{and} \quad g_0^{(TT,\,2)}(\mu) = \frac{1}{6}\,1\left(\{3 \leq |\mu| \leq 6\}\right).$$

The most favorable statistics for each choice are shown in the top row of Figure 1. Interestingly, for the thin tower model $g_0^{(TT,\,1)}$, the towers appear to be too close to each other for them to be adequately distinguishable after adding standard Gaussian noise to the clean observations μ_i; thus, the most favorable statistic



Two narrow towers

Two thick towers



Gaussian with Spike at 0

Gaussian with Spike at 2

Figure 1: Most favorable statistics for various carriers non-Gaussian carriers g_0. In each case, the noise was standard Gaussian $\varepsilon \sim \mathcal{N}(0, 1)$; the different carriers g_0 are described in Section 2.4. The relative information coefficient ρ is computed univariately for each candidate statistic.

is effectively the mean $T(\mu) = \mu$. Conversely, with $g_0^{(TT,\,2)}$, the most favorable statistic is an indicator function for the relative magnitude of each tower, and the subsequent statistics measure tilts and spreads within each tower.

We also examined the case where g_0 is a half-and-half mixture of a centered Gaussian with variance $\sigma^2 = 2$, and a δ-spike at either $\mu = 0$ or $\mu = 2$. With a spike at $\mu = 0$, the most favorable statistics are not that different from the Hermite polynomials underlying (3). However, the spike at $\mu = 2$ changes the picture considerably: now, the most stable statistic looks more like a hinge $T(\mu) = (\mu - 2)_-$.

Of course, it is not immediately clear how these non-Gaussian most favorable families are relevant to practical data analysis. A strength of the log-polynomial model (3) is that it is most favorable near *any* Gaussian carrier g_0; meanwhile, the statistics presented in Figure 1 depend on exact knowledge of g_0. However, whether or not they are immediately useful, the results presented in Figure 1 can help us gain a better feeling for the flavor of density deconvolution, and for how the shape of the carrier g_0 affects the nature of the problem.

3 Adaptive Local Minimaxity

In the previous section, we showed that there exist most favorable statistics T that capture most of the relevant information in the X-sample in terms of our relative efficiency criterion. In this section, we show how to translate this result into a more classical minimax framework. Beyond being of independent interest, the minimax results provided here may also serve as additional evidence that the relative efficiency framework proposed above is "natural" in the sense that it helps motivate good statistical procedures.

In order to describe the asymptotics of Efron's method, we need a certain amount of formalism. To this end, we focus on the local perturbation model described in Section 1.1, and show that it induces a sequence of statistical problems that converge to a locally asymptotically normal experiment [Le Cam, 1960]. This connection then lets us draw from the extensive literature on estimation in the Gaussian sequence model.

Here, we focus on the case where the carrier g_0 is Gaussian; this lets us cut down on linear algebra and to give closed-form bounds for the minimax shortfall of Efron's method. Following this choice, we use notation g_σ instead of g_0 for the carrier; here, σ^2 denotes the variance of the μ_i. We note, however, that exactly the same arguments can be used to derive the limiting risk of any most-favorable family of the type considered in Theorem 3; the minimax performance of Efron's method then depends on the decay rate of the spectrum of P_{g_0} (21). We expand on this connection in Section 3.2.

Following our discussion in Section 1.1, we are interested in a sequence of models

$$g^{(n)}(\mu) = g_\sigma^M(\mu) \exp\left[\frac{1}{\sqrt{n}}\tau(\mu) - \psi_n\right], \quad \mu \in \Omega_M = [-M,\, M], \tag{28}$$

where we eventually want to take $M \to \infty$. We assume that the tilting function $\tau : \mathbb{R} \to \mathbb{R}$ is a generic smooth function of μ contained within an ellipsoid defined with respect to its Hermite expansion

$$\tau(\mu) \in A^{\sigma}_{\kappa, C}, \quad A^{\sigma}_{\kappa, C} = \left\{ \sum_{j=1}^{\infty} \kappa^{2j} \left\langle \tau(\cdot), H_j\left(\frac{\cdot}{\sigma}\right) \right\rangle_{g_\sigma}^2 \leq C^2 \right\}. \tag{29}$$

Here, the parameters κ and C let us tune the shape of the ellipsoid, and the inner product notation is short for

$$\left\langle \tau(\cdot), H_j\left(\frac{\cdot}{\sigma}\right) \right\rangle_{g_\sigma} = \int_{\mathbb{R}} \tau(\mu) H_j\left(\frac{\mu}{\sigma}\right) g_\sigma(\mu) \, d\mu. \tag{30}$$

Our goal is to get a good estimate $\hat{g}^{(n)}$ for $g^{(n)}$ from n independent samples X_i; we measure loss in terms of the re-scaled Kullback-Leibler divergence (8). Throughout our analysis, we will assume that $\hat{g}^{(n)}$ is chosen such as to ensure that this loss is finite.

Our key result is that, given an appropriate choice of p, Efron's polynomial log-density model (3) attains quasi-minimax performance simultaneously over this class of problems, regardless of our choice of σ^2 and κ. We note that the constant on the right-hand side of (31) can be quite small. For example, if we take the carrier to be standard Gaussian $\sigma^2 = 1$, and set the ellipse shape parameter to $\kappa = 2$, we have $\alpha_{\sigma,\kappa} \approx 1.7$ and $\beta_{\sigma,\kappa} = 3$. Thus, the polynomial log-density model (3) gets to within a factor 1.8 of the minimax risk.

Theorem 6. *Suppose that we are trying to solve the sequence of problems defined by (28), (29), and (8). Let $\hat{g}^{(n,p)}$ denote maximum likelihood estimator in the polynomial log-density model (3) with p chosen as in (35), and let $\hat{g}^{(n,*)}$ be the minimax optimal estimator over the regularity class defined in (38). Then, there is a constant $C_{\sigma,k}$ such that, for $C \geq C_{\sigma,\kappa}$,*

$$\lim_{M \to \infty} \limsup_{n \to \infty} \left\{ \frac{\sup_{\tau \in A^{\sigma}_{\kappa,C}} L_n\left(\hat{g}^{(n,p)}\right)}{\sup_{\tau \in A^{\sigma}_{\kappa,C}} L_n\left(\hat{g}^{(n,*)}\right)} \right\} \leq \frac{\beta_{\sigma,\kappa}}{\alpha_{\sigma,\kappa}}, \tag{31}$$

where the constants $\alpha_{\sigma,\kappa}$ and $\beta_{\sigma,\kappa}$ are given in (33) and (37) respectively.

To establish Theorem 6, we first need to derive the worst-case risk of the minimax estimator $\hat{g}^{(n,*)}$. The key step in the proof is to show that the estimation problem outlined above converges to an elliptically constrained Gaussian sequence model. Now, as shown by Pinsker [1980], linear estimators get to within a constant factor of the minimax risk for this class of problems, and so the minimax estimation problem reduces to a linear estimation problem that can be solved directly; see Johnstone [2011] for a review. The factor 4/5 in the constant (33) is a bound obtained by Donoho et al. [1990] for the sub-optimality constant in Pinsker's result.

Lemma 7. *Under the conditions of Theorem 6, for* $C \geq C_{\sigma,\kappa}$,

$$\lim_{M\to\infty} \liminf_{n\to\infty} \left\{ \inf_{\hat{g}^{(n,*)}\in\Lambda_\sigma} \sup_{\tau\in A^\sigma_{\kappa,C}} L_n\left(\hat{g}^{(n,*)}\right) \right\} \geq \alpha_{\sigma,\kappa}\, C^{\frac{2\log(r_\sigma)}{\log(r_\sigma)+\log(\kappa)}}, \tag{32}$$

where Λ^σ_κ *is an* L_2 *regularity class defined in* (38), *and*

$$\alpha_{\sigma,\kappa} = \frac{4}{5}\frac{r_\sigma^2(\kappa-1)}{(r_\sigma^2-1)(r_\sigma^2\kappa-1)}\left(\frac{(r_\sigma^2\kappa-1)(r_\sigma^2\kappa^2-1)}{r_\sigma^2\kappa(\kappa-1)}\right)^{\frac{2\log(r_\sigma)}{\log(r_\sigma)+\log(\kappa)}}, \tag{33}$$

and r_σ *is short-hand for the constant*

$$r_\sigma^2 = \left(1+\sigma^2\right)/\sigma^2. \tag{34}$$

We can also derive the risk of $\hat{g}^{(n,p)}$ using the machinery developed for proving Lemma 7. Comparing (32) and (36), we see that both estimators have the same dependence on the signal scale C, and only differ by a constant function of σ^2 and κ.

Lemma 8. *Suppose that the true density* $g^{(n)}$ *is as in the statement of Theorem 6, and we estimate it using the exponential family model* $\hat{g}^{(n,p)}$ (3) *with*

$$p = \max\left\{2, \left\lceil \frac{\log(C/\sigma)}{\log(r_\sigma\kappa)} \right\rceil - 1\right\}. \tag{35}$$

Then, for $C \geq C_{\sigma,\kappa}$,

$$\lim_{M\to\infty} \limsup_{n\to\infty} \left\{ \sup_{\tau\in A^\sigma_{\kappa,C}} L_n\left(\hat{g}^{(n,p)}\right) \right\} \leq \beta_{\sigma,\kappa} C^{\frac{2\log(r_\sigma)}{\log(r_\sigma)+\log(\kappa)}}, \tag{36}$$

with

$$\beta_{\sigma,\kappa} = \left(1+r_\sigma^2\right)\sigma^{\frac{2\log(\kappa)}{\log(r_\sigma)+\log(\kappa)}}. \tag{37}$$

3.1 Density Deconvolution and the Gaussian Sequence Model

Here, we outline the proof of Lemma 7 by showing how our density estimation problem converges to a Gaussian sequence model. Throughout our analysis, we assume that our density estimate is in the class Λ^σ_κ defined by the following relation

$$\hat{g}^*_n \in \Lambda^\sigma_\kappa, \quad \Lambda^\sigma_\kappa = \left\{ g : \sum_{j=1}^{\infty} \kappa^{2j} \left\langle \log\left(\frac{g(\cdot)}{g^M_\sigma(\cdot)}\right), H_j\left(\frac{\cdot}{\sigma}\right)\right\rangle_{g^M_\sigma} < L^2 \right\} \tag{38}$$

for some large constant $L^2 > C^2$, along with the constraint that all integrals defined above be finite and well-defined.

Thanks to our regularity assumption, we can use notation from (30) to define "Hermite coefficients"

$$\gamma_j = \left\langle \tau(\cdot), H_j\left(\frac{\cdot}{\sigma}\right)\right\rangle_{g_\sigma}. \tag{39}$$

Meanwhile, given our sample X_1, ..., X_n, we can also define empirical Hermite coefficients as

$$Z_j^{(n)} = \frac{1}{\sqrt{n}}\left(\frac{1+\sigma^2}{\sigma^2}\right)^{j/2}\sum_{i=1}^{n} H_j\left(\frac{X_i}{\sqrt{1+\sigma^2}}\right). \tag{40}$$

Because $g^{(n)}$ converges to the carrier g_0 as n gets large, we can show that for any finite set of indices $\{j_i\}$, the $Z_{j_i}^{(n)}$ are asymptotically jointly normal with

$$\lim_{n\to\infty}\mathbb{E}\left[Z_{j_i}^{(n)}\right] = \gamma_{j_i} + o_M(1), \tag{41}$$

$$\lim_{n\to\infty}\mathrm{Cov}\left[Z_{j_i}^{(n)}, Z_{j_{i'}}^{(n)}\right] = \delta_{\{j_i=j_{i'}\}} r_\sigma^{2j_i} + o_M(1), \tag{42}$$

where we again used the notation $r_\sigma^2 = \left(1+\sigma^2\right)/\sigma^2$.

These observations suggest that deriving a good estimator $\hat{g}^{(n)}$ for the density $g^{(n)}$ is related to finding the mean of the Gaussian sequence with covariance structure (41). The following lemma makes this connection explicit, thus reducing our setting to a well-understood problem, namely estimating a Gaussian sequence model with elliptical constraints under squared-error loss.

Lemma 9. *For large M, the limiting minimax risk for the sequence of density deconvolution problems defined in Lemma 7 converges to the minimax risk for estimating the mean of a Gaussian sequence Z_j for $j = 1, 2, \ldots$ with*

$$\mathbb{E}[Z_j] = \gamma_j,\ \mathrm{Cov}[Z_j, Z_k] = \delta_{\{j=k\}} r_\sigma^{2j} \tag{43}$$

under the loss

$$L\left(\hat{\gamma}(Z)\right) = \sum_{j=1}^{\infty}\left(\hat{\gamma}(Z) - \gamma\right)^2, \tag{44}$$

and the constraint $\gamma \in \ell_2^\kappa(C) := \left\{\gamma : \sum_{j=1}^{\infty}\kappa^{2j}\gamma_j^2 \le C^2\right\}$. In other words,

$$\lim_{M\to\infty}\lim_{n\to\infty}\left\{\inf_{\hat{g}\in\Lambda_\kappa^\sigma}\sup_{\tau\in A_{\kappa,C}^\sigma} L_n\left(\hat{g}^{(n)}\left(Z^{(n)}\right)\right)\right\} = \inf_{\hat{\gamma}\in\ell_2^\kappa(C)}\sup_{\gamma\in\ell_2^\kappa(C)} L\left(\hat{\gamma}(Z)\right). \tag{45}$$

Now, this class of Gaussian sequence models can be well estimated using linear rules: Pinsker [1980] established that the risk of the best linear rule is within a constant factor of the minimax risk; furthermore, Donoho et al. [1990] showed that this constant is less than 5/4. Thus, it suffices to find the risk of the best linear rule of the form $\hat{\gamma}_j^L = c_j Z_j$ for some constants c_j. Now, in the

Gaussian limit, the worst-case risk of the minimax linear rule $\hat{\gamma}^L$ has a simple form [e.g., Johnstone, 2011, Chapter 5.1]:

$$R^L := \sup_{\gamma \in \ell_2^s(C)} \mathbb{E}\left[L\left(\hat{\gamma}^L\right)\right] = \sum_{j=1}^{\infty} r_\sigma^{2j} \left(1 - \frac{\kappa^j}{\mu_C}\right)_+, \tag{46}$$

where μ_C is implicitly defined by the relation

$$\sum_{j=1}^{\infty} r_\sigma^{2j} \kappa^j \left(\mu_C - \kappa^j\right)_+ = C^2. \tag{47}$$

Since we know that the minimax risk of R^* is bounded from below by $4R^L/5$, the proof of Lemma 7 reduces to algebra; the remaining steps are carried out in the appendix.

3.2 Extension to General Carriers

In an effort to reduce notational burden, the above argument focused on the Gaussian carrier case, i.e. $g_0 = g_\sigma$ for some $\sigma > 0$. In this section, we briefly outline the technical ideas needed to prove analogous local minimaxity results in the neighborhood of general carriers g_0. The proof of Theorem 6 relied on showing that large-sample statistical inference in the local density deconvolution model reduces to a study of the asymptotically normal "empirical Hermite coefficients" $Z_j^{(n)}$ defined in (40). A similar result also holds in the case of general g_0; however, the statistics $Z_j^{(n)}$ now take on a more general form.

Let $\{T_j(\mu)\}_{j=1}^{\infty}$ be the statistics comprising the most-favorable family around g_0, and denote the relative efficiency of the j-dimensional most-favorable family by ρ_j. Then, inference can be asymptotically framed in terms of the moments

$$Z_j^{(n)} = \frac{1}{\sqrt{n}} \rho_j^{-1} \sum_{i=1}^{n} U(X_i), \quad \text{with} \quad U(x) := \frac{\left(\phi * \left(T_j\, g_0\right)\right)(x)}{f_0\,(x)}. \tag{48}$$

The following lemma shows that the $Z_j^{(n)}$ in fact have the desired limiting distribution. An analogue to Lemma 9 then follows directly.

Lemma 10. *Suppose that the conditions of Theorem 3, and that the perturbation function τ satisfies $\int_\Omega \tau^2\,(\mu)\, g_0\,(\mu)\; d\mu < \infty$. Then, then the $Z_j^{(n)}$ defined in (48) are asymptotically normal over any finite set of indices and, for any j, $j' \in \mathbb{N}$,*

$$\lim_{n\to\infty} \mathbb{E}\left[Z_j^{(n)}\right] = \int_\Omega T_j\,(\mu)\; \tau\,(\mu)\; g_0\,(\mu)\; d\mu, \tag{49}$$

$$\lim_{n\to\infty} \mathrm{Cov}\left[Z_j^{(n)},\, Z_{j'}^{(n)}\right] = \delta_{\{j=j'\}}\, \rho_j^{-1}. \tag{50}$$

Given this result, we can use Pinsker's theorem to compute tight bounds for the minimax risk of density deconvolution near g_0 as a function of the most-favorable relative efficiency coefficients ρ_j (i.e., the spectrum of P_{g_0} defined in (21)). We note that Lemma 10—and in fact the whole machinery of using a Gaussian sequence model to understand the behavior of maximum likelihood estimation in the most-favorable family—only depends on the square-integrability condition $\int_\Omega \tau^2(\mu)\, g_0(\mu)\ d\mu < \infty$ and on the compactness of P_{g_0}. However, in order for the resulting minimaxity properties to be any good, we need for the spectrum of P_{g_0} to decay fast.

4 Two Optimality Theories for Density Deconvolution

The main contribution of our paper is a local optimality theory for density deconvolution that helps us understand and justify Efron's method, i.e., maximum likelihood estimation in a model of the form (2). This line of work is in contrast to the classical optimality theory based on kernel estimators. Notable contributions include the pioneering work of Carroll and Hall [1988] and Stefanski and Carroll [1990], the analysis of Fan [1991] that elucidates the connection between the decay rate of the Gaussian characteristic function and the difficulty of density deconvolution, the strong quasi-minimaxity results of Efromovich [1997], as well as several other papers cited in the introduction.

A key difference between these theories is that kernel-based methods are quasi-optimal in terms of an integrated squared error (ISE) loss criterion whereas our theory is framed in terms of the Kullback-Leibler (KL) loss, defined respectively as

$$\mathrm{ISE}(g,\, \hat{g}) = \int (\hat{g}(\mu) - g(\mu))^2\, d\mu, \tag{51}$$

$$\mathrm{KL}(g,\, \hat{g}) = \int \log\left(\frac{g(\mu)}{\hat{g}(\mu)}\right) g(\mu)\, d\mu. \tag{52}$$

Moreover, our theory aims to detect *local* perturbations of g_0, whereas the ISE criterion is usually applied *globally*.

From a scientific point of view, the value of an optimality theory depends on the relevance of the induced estimators to answering real-world questions. Kernel-based methods have a good track record for solving some classic problems, such as in-season baseball prediction problem introduced by Efron and Morris [1975]; see, e.g., Brown [2008]. In fact, these methods have been shown to approach the Bayes risk for estimating the posterior mean $\mathbb{E}\left[\mu \,\middle|\, X = x\right]$ [Brown and Greenshtein, 2009].

In this section, however, we present examples of natural scientific questions for which Efron's method provides substantially better answers than kernel methods. We hope that these examples will convince the reader that a KL-based optimality theory for density deconvolution is, if nothing else, worthy of

further study. For our experiments, we used the `R`-package `decon` [Wang and Wang, 2011] for kernel-based estimation.

4.1 Detecting the Fraction of Weakly Associated Genes in a Microarray Study

Our first example simulates a classic biological application: gene expression profiling using microarrays. At a high level, the goal is to estimate the difference in the expression levels of different genes for two distinct cell populations (e.g., breast cancer cells vs. healthy cells). After some pre-processing, each gene can be associated with a test statistic that has a standard normal distribution under the null hypothesis that the gene's expression levels do not differ accross the two groups; the resulting statistical problem is described in detail by Efron et al. [2001] and Tusher et al. [2001].

Here, we follow the *structural model* analysis of Efron [2004] who showed that, to reasonable approximation, we can model the gene-wise test statistics X_i as $X_i \sim \mathcal{N}(\mu_i, 1)$, where μ_i describes the true association between the i-th gene and the condition of interest. The i-th "exact" null hypothesis is that $\mu_i = 0$. However, as argued in, e.g., Chapter 6 of Efron [2010], this exact null may not always be scientifically relevant. It seems likely that most genes have a small but non-zero true association μ_i; the goal of the statistician is then not to identify the non-zero μ_i but rather to identify the *large* associations μ_i.

Motivated by this setup, suppose that—given the power afforded by our sample size—we decide that having $\mu_i \in [-2, 2]$ qualifies as a "small" association, and we want to estimate the fraction of genes whose association is in this range. In our framework, assuming that the gene associations μ_i are drawn from a distribution with density $g(\cdot)$, we want to estimate $\int_{-2}^{2} g(\mu)\, d\mu$. In Figure 2, we ran 1,000 replicates of such a simulation with $n = 5,000$ genes each, where the true density g was given by

$$g(\mu) = 0.95 \cdot \frac{1}{4}(2 - |\mu|)_+ + 0.05 \cdot \frac{1}{20} 1(\{|\mu| \leq 10\}).$$

We then generated X-statistics from the structural model $X \sim \mathcal{N}(\mu, 1)$, and generated estimates $\hat{g}$ both using the kernel method `decon` and maximum likelihood in the family (3) with $p = 4$.

In this simulation, the correct answer was $\int_{-2}^{2} g(\mu)\, d\mu = 0.96$. Efron's method (3) does substantially better than the kernel method, with the exception of a few cases where it vastly underestimates the bulk of $g(\cdot)$. In practice, Efron [2014a] recommends the use of regularization to stabilize the estimators and mitigate such finite sample effects. Here, however, we did not use any regularization in order to keep our experiments as simple as possible.

4.2 Identifying Safe Neighborhoods

Our second example is based on a real dataset: the communities and crime unnormalized dataset from the UCI Machine Learning Repository [Lichman, 2013],

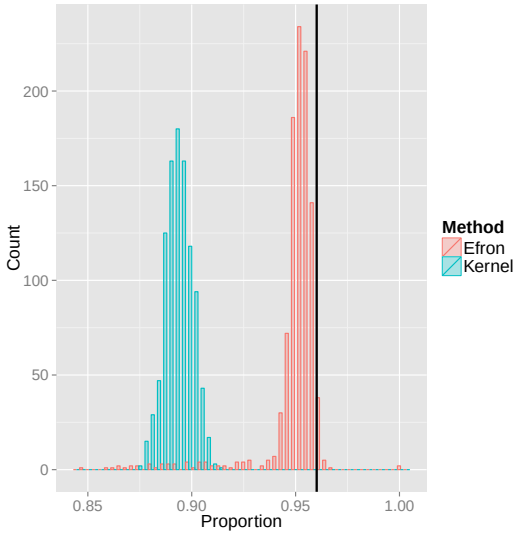


Figure 2: Gene expression profiling simulation. The goal is to estimate $\int_{-2}^{2} g(\mu)\, d\mu$; here, the correct answer (0.96) is indicated by a thick vertical line. Over 1,000 simulation replicates, Efron's method usually performs substantially better than the kernel-based alternative.

which tallies crime data from multiple US communities in 1995; we restricted our analysis to the $n = 1,192$ communities with population greater than 20,000 and non-missing crime data.

Our goal was to understand the number of non-violent crimes per 100 people. In our sample of communities, the average non-violent crime rate was around 5 per 100 inhabitants, while 6% of communities achieved a non-violent crime rate of 2%. We define a community with a rate of less than 2% as *safe*. The dataset was large enough that we could accurately estimate per-community crime rates; to test our deconvolution methods, we made the problem harder by down-sampling the data and then seeking to recover the correct answer.

To down-sample the data in a mathematically principled way, we made the practically somewhat implausible assumption that each person was victim of at most 1 crime in 1995. Then, we can imagine assembling a crime dataset by interviewing $B = 500$ randomly selected people per community, and counting the number N_i of interviewees in the i-th community that have been victims of a non-violent crime. We can easily simulate the outcome of such an interview using the available data by hypergeometric sampling:

$$N_i \sim \text{Hyper}\left(B, \text{ Crimes in community } i, \text{ Population } i\right).$$

Our statistical task is to estimate which communities are *safe* (i.e., have a rate of less than 2%) based on statistics N_i collected in different communities. Let p_i denote the true crime rate in the i-th community, and let $\hat{p}_i = N_i/B$. By a

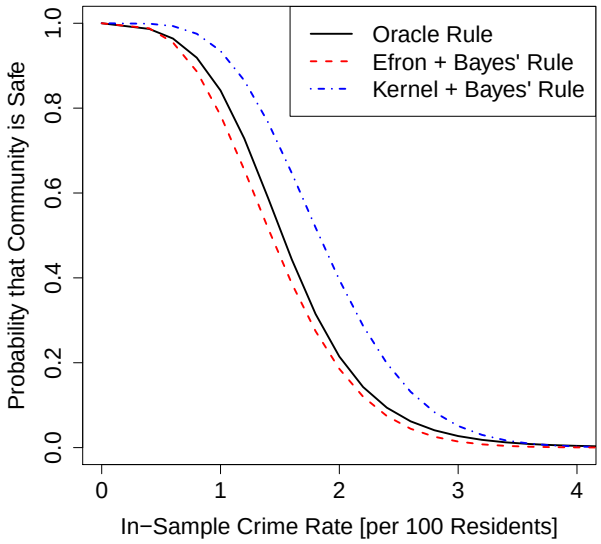


Figure 3: Crime prediction example. The goal is to predict the probability that a community is safe ($p_i \leq 0.02$) given a crime-rate estimate $\hat{p}_i$ obtained by interviewing $B = 500$ randomly selected people. Efron's method provides a closer approximation of the oracle rule than the kernel method.

standard variance stabilizing argument, we can verify that

$$\sqrt{\hat{p}_i} \dot{\sim} \mathcal{N}\left(\sqrt{p_i}, \frac{1}{B}\right).$$

Thus, to estimate $\mathbb{P}\left[p_i \leq 0.02 | \hat{p}_i\right]$, we can first apply a square-root transform to the data, then use any method to estimate the density g of the $\sqrt{p_i}$, and finally apply Bayes' rule to get our desired quantity.

Figure 3 shows results for the kernel method and Efron's method; both density deconvolution methods were implemented exactly the same way as for the gene expression example. The oracle rule was produced by fitting a spline logistic regression of $1\left(\{p_i \leq 0.02\}\right)$ against $\hat{p}_i$. Efron's method again substantially outerforms the alternative. For example, if we ask about the probability that a community is safe given that $\hat{p}_i = 0.02$, Efron's method tells us that this probability is 18.6% whereas the kernel method tells us 39.6%. The correct answer, produced by the oracle fit, is 21.5%.

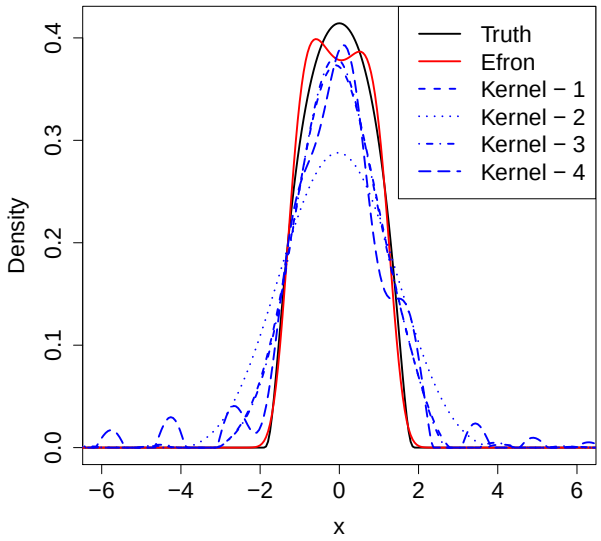


Figure 4: A comparison of Efron's method with a log-polynomial model (3) with $p = 4$ with kernel-based density deconvolution as implemented in `decon`. We tried four different bandwidth choices for the kernel method.

4.3 Fitting the Mode vs. Fitting the Support

To provide some deeper insight into the differences between the two density deconvolution methods under consideration, we end our experimental section with an in-depth look at a simple density estimation problem with a large sample size: we drew $n = 100,000$ observations from the generative model (1) with a density g of the form

$$g(\mu) \propto \exp\left[-1/\left(1 + \frac{4}{\mu^2}\right)\right] \text{ for } -2 < \mu < 2, \text{ and } g(\mu) = 0 \text{ else.}$$

We note that this density is not contained in the span of either (3) or the basis functions implicitly used by `decon`. Results are shown in Figure 4.

To give the kernel density method a good chance of performing well, we tried four different bandwidth-selection algorithms provided by `decon`: **(1)** a closed form approximation to the boostrap recommended by Delaigle and Gijbels [2004] (this method is default, and was used for the other experiments), **(2)** a rule of thumb by Fan [1991], **(3)** an approximate ISE minimzer by Stefanski and Carroll [1990], and finally **(4)** a hand-selected bandwidth to demonstrate the effect of mild under-smoothing. For Efron's method we, as usual, set $p = 4$ in (3).

We notice that all methods do a roughly equivalent job of estimating the

target density g near its mode. However, Efron's method is much more accurate in estimating the tails and support of g. This observation is not entirely surprising in light of the loss functions that were used to motivate each method. The ISE loss (51) only requires $\hat{g}$ to be reasonable close to g over the domain of μ, but does not impose particularly harsh penalties for oscillatory behavior in the tails. In contrast, the KL loss places more attention on getting the tails and support of the distribution right. It thus appears that kernel methods built on ISE-optimality theory are unreliable for answering scientific questions that depend on understanding the tail-behavior of g, whereas methods based on KL-optimality may perform better.

5 Discussion

The work of Efron [2014a,b] presents a surprising challenge to the theory of density deconvolution. From, e.g., Efromovich [1997], it may appear that the problem of density deconvolution with Gaussian errors is completely solved—and that kernel density estimators are optimal for the task. And yet, Efron [2014b] found that his maximum likelihood method vastly out-performed methods related to kernel density estimators for several realistic problems (Efron calls this latter approach "f-estimation").

The goal of this paper was to understand why Efron's method could work better than kernel density estimators despite the theoretical guarantees available for the latter. To do so, we introduced a perturbation model inspired by Le Cam's local asymptotic normality theory, and showed that Efron's method is quasi-optimal in this setup for deviance loss. Under the assumption that deviance loss comes closer to describing the "real loss function" of a practitioner than the integrated squared error loss used to establish optimality properties of kernel density estimators, our results can be seen as moving towards an explanation for the empirical success of Efron's method.

More broadly, our results highlight a surprisingly strong connection between non-parametric density deconvolution and classical likelihood theory. We found that—at least locally—there are some directions along which we can accurately estimate a signal and maximum-likelihood estimation is quasi-optimal for this task; meanwhile, there are other directions in which estimation is hopeless and the minimax strategy is to ignore them. From a practical perspective, this connection appears rather reassuring, as low-dimensional large-sample maximum-likelihood estimation has proven to be one of the most consistently successful ideas in applied statistics.

References

L D Brown. In-season prediction of batting averages: A field test of empirical Bayes and Bayes methodologies. *The Annals of Applied Statistics*, 2(1):113–152, 2008.

L D Brown and E Greenshtein. Nonparametric empirical Bayes and compound decision approaches to estimation of a high-dimensional vector of normal means. *The Annals of Statistics*, 37(4):1685–1704, 2009.

L D Brown and M G Low. Asymptotic equivalence of nonparametric regression and white noise. *The Annals of Statistics*, 24(6):2384–2398, 1996.

L D Brown, A V Carter, M G Low, and C H Zhang. Equivalence theory for density estimation, Poisson processes and Gaussian white noise with drift. *The Annals of Statistics*, 32(5):2074–2097, 2004.

C Butucea and F Comte. Adaptive estimation of linear functionals in the convolution model and applications. *Bernoulli*, 15(1):69–98, 2009.

R J Carroll and P Hall. Optimal rates of convergence for deconvolving a density. *Journal of the American Statistical Association*, 83(404):1184–1186, 1988.

R J Carroll and P Hall. Low order approximations in deconvolution and regression with errors in variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(1):31–46, 2004.

F Comte and C Lacour. Data-driven density estimation in the presence of additive noise with unknown distribution. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):601–627, 2011.

A Delaigle and I Gijbels. Bootstrap bandwidth selection in kernel density estimation from a contaminated sample. *Annals of the Institute of Statistical Mathematics*, 56(1):19–47, 2004.

D L Donoho, R C Liu, and B MacGibbon. Minimax risk over hyperrectangles, and implications. *The Annals of Statistics*, 18(3):1416–1437, 1990.

S Efromovich. Density estimation for the case of supersmooth measurement error. *Journal of the American Statistical Association*, 92(438):526–535, 1997.

S Efromovich and A Samarov. Asymptotic equivalence of nonparametric regression and white noise model has its limits. *Statistics & Probability Letters*, 28 (2):143–145, 1996.

B Efron. Large-scale simultaneous hypothesis testing. *Journal of the American Statistical Association*, 99(465):96–104, 2004.

B Efron. *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Cambridge University Press, 2010.

B Efron. The Bayes deconvolution problem. 2014a.

B Efron. Two modeling strategies for empirical Bayes estimation. *Statistical Science*, 29(2):285, 2014b.

B Efron and C Morris. Data analysis using Stein's estimator and its generalizations. *Journal of the American Statistical Association*, 70(350):311–319, 1975.

B Efron, R Tibshirani, J D Storey, and V Tusher. Empirical Bayes analysis of a microarray experiment. *Journal of the American Statistical Association*, 96 (456):1151–1160, 2001.

S Ehrenfeld. On the efficiency of experimental designs. *The Annals of Mathematical Statistics*, 26(2):247–255, 1955.

J Fan. On the optimal rates of convergence for nonparametric deconvolution problems. *The Annals of Statistics*, 19(3):1257–1272, 1991.

J Fan and J Y Koo. Wavelet deconvolution. *IEEE Transactions on Information Theory*, 48(3):734–747, 2002.

G K Golubev, M Nussbaum, and H H Zhou. Asymptotic equivalence of spectral density estimation and Gaussian white noise. *The Annals of Statistics*, 38(1): 181–214, 2010.

P Hall and S N Lahiri. Estimation of distributions, moments and quantiles in deconvolution problems. *The Annals of Statistics*, 36(5):2110–2134, 2008.

P Hall and A Meister. A ridge-parameter approach to deconvolution. *The Annals of Statistics*, 35(4):1535–1558, 2007.

W Jiang and C H Zhang. General maximum likelihood empirical Bayes estimation of normal means. *The Annals of Statistics*, 37(4):1647–1684, 2009.

I M Johnstone. Gaussian estimation: Sequence and wavelet models. 2011.

R Koenker and I Mizera. Convex optimization, shape constraints, compound decisions, and empirical Bayes rules. *Journal of the American Statistical Association*, 109(506):674–685, 2014.

J Y Koo. Logspline deconvolution in Besov space. *Scandinavian Journal of Statistics*, 26(1):73–86, 1999.

J Y Koo and H Y Chung. Log-density estimation in linear inverse problems. *The Annals of Statistics*, 26(1):335–362, 1998.

L Le Cam. *Locally asymptotically normal families of distributions*, volume 3, pages 27–98. University of California Publ. in Statistics, 1960.

M Lichman. UCI machine learning repository, 2013. URL http://archive.ics.uci.edu/ml.

T A Louis. Finding the observed information matrix when using the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 44 (2):226–233, 1982.

M Nussbaum. Asymptotic equivalence of density estimation and Gaussian white noise. *The Annals of Statistics*, 24(6):2399–2430, 1996.

M S Pinsker. Optimal filtering of square-integrable signals in Gaussian noise. *Problemy Peredachi Informatsii*, 16(2):52–68, 1980.

J Shawe-Taylor, C Williams, N Cristianini, and J Kandola. On the eigenspectrum of the Gram matrix and the generalization error of kernel-PCA. *IEEE Transactions on Information Theory*, 51(7):2510–2522, 2005.

L A Stefanski and R J Carroll. Deconvoluting kernel density estimators. *Statistics: A Journal of Theoretical and Applied Statistics*, 21(2):169–184, 1990.

V G Tusher, R Tibshirani, and G Chu. Significance analysis of microarrays applied to the ionizing radiation response. *Proceedings of the National Academy of Sciences*, 98(9):5116–5121, 2001.

A W van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 2000.

A Wald. Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, 46(2):265–280, 1945.

M P Wand. Finite sample performance of deconvolving density estimators. *Statistics & Probability Letters*, 37(2):131–139, 1998.

X F Wang and B Wang. Deconvolution estimation in measurement error models: The R package decon. *Journal of Statistical Software*, 39(10), 2011.

C H Zhang. Fourier methods for estimating mixing densities and distributions. *The Annals of Statistics*, 18(2):806–831, 1990.

C H Zhang. Generalized maximum likelihood estimation of normal mixture densities. *Statistica Sinica*, 19(3):1297, 2009.

A Proofs

Proof of Lemma 1

It suffices to prove the second conclusion, as the alternative expression for the relative efficiency coefficient given in (14) can directly be verified to be transformation invariant. We already know that $\mathcal{I}_\mu(\eta) = \mathrm{Var}_\eta[t(\mu)]$, and so all we really need to do is to check that

$$\mathcal{I}_X(\eta) = \mathrm{Var}_\eta\left[\mathbb{E}_\eta\left[t(\mu) \,\middle|\, X\right]\right].$$

Although we could also get to this answer by drawing analogies to the work of Louis [1982], we will take a direct approach here. In the univariate family with statistic t, the expected score function is

$$\begin{aligned}
\frac{d}{d\eta}\log f_\eta(x) &= \frac{d}{d\eta}\log\int K(\mu, x)\, g_\eta(\mu)\, d\mu \\
&= \frac{\int K(\mu, x)\, g'_\eta(\mu)\, d\mu}{\int K(\mu, x)\, g_\eta(\mu)\, d\mu} \\
&= \frac{\int K(\mu, x)\, g_\eta(\mu)\,(t(\mu) - \mathbb{E}_\eta[t(\mu)])\, d\mu}{\int K(\mu, x)\, g_\eta(\mu)\, d\mu} \\
&= \mathbb{E}_\eta\left[t(\mu) \,\middle|\, X = x\right] - \mathbb{E}_\eta[t(\mu)].
\end{aligned}$$

Thus, we conclude that

$$\begin{aligned}
\mathcal{I}_X(\eta) &= \mathrm{Var}_\eta\left[\frac{d}{d\eta}\log f_\eta(x)\right] \\
&= \mathrm{Var}_\eta\left[\mathbb{E}_\eta\left[t(\mu) \,\middle|\, X = x\right]\right].
\end{aligned}$$

Proof of Lemma 2

By Bayes' rule, we know that the conditional density of μ given x is

$$g\left(t(\mu) \,\middle|\, X = x\right) = \frac{K(x, \mu)\, g(\mu)}{f(x)};$$

from this, it directly follows that

$$\mathrm{Var}\left[\mathbb{E}\left[t(\mu) \,\middle|\, X\right]\right] = \int_\Omega t^2(\mu)\, K^2(x, \mu)\, g^2(\mu)\, f^{-1}(x)\, dx\, d\mu.$$

Similarly, we can check that

$$\mathrm{Var}\left[t\left(\mu\right)\right] = \int_{\Omega} t^2\left(\mu\right) g\left(\mu\right) d\mu.$$

For a statistic $t\left(\mu\right)$ of the form $t = a \cdot T$, we can write $t^2 = a^\top (TT^\top) a$; with this notation, we recover the first result of Lemma 2. The explicit solution given in (19) and (20) for a^* in terms of the spectrum of $Q_T^\top M_T Q_T$ is standard.

Proof of Theorem 3

Because the relative efficiency coefficient is transformation invariant, we can without loss of generality pick a statistic T for which

$$\int_{\Omega} T\left(\mu\right) g_0\left(\mu\right) d\mu = 0, \tag{53}$$

$$\int_{\Omega} T\left(\mu\right) T^\top(\mu)\, g_0\left(\mu\right) d\mu = I_{p\times p}. \tag{54}$$

Given such a choice, the relative efficiency formula from Lemma 2 simplifies to

$$\begin{aligned}\rho_0\left(T\right) &= \lambda_{\min}\left(\int_{\Omega} T\left(\mu\right) T^\top(\mu)\, K^2\left(x,\, \mu\right) g_0^2\left(\mu\right) f^{-1}\left(x\right) dx\, d\mu\right), \\ &= \lambda_{\min}\left(\int_{\Omega} \left(\sqrt{g_0}\left(\mu\right) T\left(\mu\right)\right) P_{g_0}\left(\mu,\, \mu\right) \left(\sqrt{g_0}\left(\mu\right) T\left(\mu\right)\right)^\top d\mu\right),\end{aligned} \tag{55}$$

where $\lambda_{\min}(A)$ denotes the smallest eigenvalue of a linear operator A. Minimizing the objective (55) subject to the constraint (54) is a standard spectral analysis problem. By the Courant-Fischer-Weyl maximin theorem as stated in, e.g., Shawe-Taylor et al. [2005], we find that because P_{g_0} is both self-adjoint and compact (and thus also completely continuous),

$$\max\left\{\rho_0\left(S\right) : S \in L_2\left(\Omega\right)^p,\, \int_{\Omega} S\left(\mu\right) S^\top(\mu)\, g_0\left(\mu\right) d\mu = I_{p\times p}\right\} = \lambda_p,$$

where λ_p is the p-the eigenvalue of P_{g_0}; moreover, this maximum is attained by setting $S_j\left(\mu\right) = \zeta_j\left(\mu\right)/\sqrt{g_0(\mu)}$ for $j = 1, ..., p$, where the ζ_j are the leading eigenvectors of P_{g_0}. We note that $\zeta_j\left(\mu\right) \in L_2\left(\Omega\right)$ because P_{g_0} is compact; S_j is then also in $L_2\left(\Omega\right)$ because g_0 is bounded away from 0 on Ω. Finally, because g_0 and K are continuous, $P_{g_0} S_j$ is continuous and so S_j must also be continuous.

Now, we still need to deal with the constraint (53). Thankfully, we can verify that all the eigenvalues of P_{g_0} are bounded by 1, and that $\zeta_1\left(\mu\right) := \sqrt{g_0\left(\mu\right)}$ is an eigenfunction of P_{g_0} with eigenvalue 1. By orthogonality of the spectrum of P_{g_0}, we then see that

$$\int_{\Omega} \frac{\zeta_j\left(\mu\right)}{\sqrt{g_0\left(\mu\right)}} g_0\left(\mu\right) d\mu = \int_{\Omega} \zeta_j\left(\mu\right) \zeta_1\left(\mu\right) = 0$$

for all $j > 1$. Thus, the minimizer of (55) with both constraints (53) and (54) is given by $T_j(\mu) = \zeta_{j+1}(\mu)/\sqrt{g_0(\mu)}$ for $j = 1, ..., p$, and the objective value (55) is the $(p+1)$-st eigenvalue of P_{g_0}. Moreover, if the spectrum of P_{g_0} does not have repeated eigenvalues, the span of $T_1, ..., T_p$ maximizing our objective is unique. We note that, because g_0 and S_j are continuous and Ω is compact, $\mathbb{E}[g_0(\mu)\exp(\eta \cdot T(\mu))]$ is finite for η in a neighborhood of 0 and so our most-stable family is in fact well-defined.

Proof of Corollary 4

Because P_{g_0} is compact, we know by the eigenfunctions $\{\zeta_j\}_{j=1}^{\infty}$ form a complete orthonormal basis for $L_2(\Omega)$; thus, we see that the statistics $\{T_j\}_{j=0}^{\infty}$ also form a complete orthonormal basis for $L_2(\Omega)$ with inner product weighted by g_0 and can write[2]

$$\tau(\mu) = \sum_{j=1}^{\infty} \gamma_j T_j(\mu), \quad \int_{\Omega} \tau^2(\mu) g_0(\mu) = \sum_{j=1}^{\infty} \gamma_j^2.$$

Given this notation, we set $\tau^{(p,*)} = \sum_{j=1}^{p} \gamma_j T_j(\mu)$. Our goal is to show that, for any integer $J > p$, our choice of $\tau^{(p,*)}$ satisfies the conclusion of Corollary 4 under the assumption that $\gamma_j = 0$ for all $j > J$. Because this bound holds uniformly in J, we conclude that it also holds in the non-parametric case $\int_{\Omega} \tau^2(\mu) g_0(\mu) \leq C^2$.

Following the above discussion, we now assume that $\tau(\mu) = \sum_{j=1}^{J} \gamma_j T_j(\mu)$, and write γ for the parameter vector inducing τ. The target loss is then

$$\begin{aligned} &D_{KL}\left(f_{\tau}^{(n)}, f_{\tau^{(p,*)}}^{(n)}\right) \\ &\quad = \int_{\Omega} f_{\tau/\sqrt{n}}(x) \log\left(\frac{f_{\tau/\sqrt{n}}(x)}{f_{\tau^{(p,*)}/\sqrt{n}}(x)}\right) dx \\ &\quad = \frac{1}{2n} \int_{\Omega} \left(\gamma^{(p,*)} - \gamma\right)^{\top} \nabla^2 \log\left(f_{\tau/\sqrt{n}}(x)\right) \left(\gamma^{(p,*)} - \gamma\right) f_{\tau/\sqrt{n}}(x)\, dx \\ &\qquad\qquad + o\left(n^{-1}\right), \end{aligned}$$

where the derivative ∇^2 is taken with respect to the parameters γ. We note that the first-order term depending on $\nabla \log(\cdot)$ integrates out to 0. Now, taking limits, we find that

$$\begin{aligned} &\lim_{n\to\infty} n\, D_{KL}\left(f_{\tau}^{(n)}, f_{\tau^{(p,*)}}^{(n)}\right) \\ &\quad = \frac{1}{2}\left(\gamma^{(p,*)} - \gamma\right)^{\top} \int_{\Omega} \nabla^2 \log\left(f_0(\mu)\right) f_0(x)\, dx \, \left(\gamma^{(p,*)} - \gamma\right) \\ &\quad = \frac{1}{2}\left(\gamma^{(p,*)} - \gamma\right)^{\top} \mathcal{I}_X^{\Gamma_J(g_0)} \left(\gamma^{(p,*)} - \gamma\right), \end{aligned}$$

[2] Without loss of generality, $\gamma_0 = 0$ since this term is absorbed by the normalization.

where $\mathcal{I}_X^{\Gamma_J(g_0)}$ denotes the Fisher information for estimating γ from X-samples in the J-dimensional most-favorable family. But now, $\Gamma_J(g_0)$ is scaled such that $\mathcal{I}_\mu^{\Gamma_J(g_0)} = I_{J\times J}$. Thus, the limiting approximation error is equal to

$$\lim_{n\to\infty} n\, D_{KL}\left(f_\tau^{(n)},\, f_{\tau^{(p,\,*)}}^{(n)}\right) = \frac{1}{2}\left\|\gamma^{(p,\,*)} - \gamma\right\|_2^2 \rho\left(\left(\gamma^{(p,\,*)} - \gamma\right)\cdot \Gamma_j(g_0)\right),$$

which, by Theorem 3, can be bounded above by $\frac{1}{2}C^2\lambda_{p+2}(P_{g_0})$.

Proof of Theorem 5

Given our assumptions, we know that the carrier g and the marginal density of the observations f are given by

$$g(\mu) = \frac{1}{\sigma}\varphi\left(\frac{\mu}{\sigma}\right), \quad f(x) = \frac{1}{\sqrt{1+\sigma^2}}\varphi\left(\frac{x}{\sqrt{1+\sigma^2}}\right),$$

where both densities loop around if needed to accommodate the bounded domain. For our proof, we begin by verifying that the

$$\nu_j(\mu) = \sqrt{\frac{1}{\sigma}\varphi\left(\frac{\mu}{\sigma}\right)}\, H_j\left(\frac{\mu}{\sigma}\right)$$

are eigenfunctions of P_g in the limit $M = \infty$. Then, for large but finite M, the ν_j are nearly eigenfunctions of P_g; meanwhile, the conditions of Theorem 3 are satistified and so our desired conclusion follows.

Now, to verify that the ν_j are eigenfunctions with $M = \infty$, we first note that P_g is a compact kernel and so it does in fact admit a spectral decomposition, and second that

$$\begin{aligned}
&\int \nu_j(\mu_1)\, P_g(\mu_1,\, \mu_2)\, \nu_k(\mu_2)\, d\mu_1 d\mu_2 \\
&\quad = \int \left(\varphi * \left(\frac{1}{\sigma}\varphi\left(\frac{\cdot}{\sigma}\right)\cdot H_j\left(\frac{\cdot}{\sigma}\right)\right)\right)(x) \\
&\qquad\qquad \left(\varphi * \left(\frac{1}{\sigma}\varphi\left(\frac{\cdot}{\sigma}\right)\cdot H_k\left(\frac{\cdot}{\sigma}\right)\right)\right)(x)\, \frac{dx}{f(x)}
\end{aligned}$$

Now, focusing on the inner terms, we can check that

$$\frac{1}{\sigma}\varphi\left(\frac{\mu}{\sigma}\right)\cdot H_j\left(\frac{\mu}{\sigma}\right) = \frac{1}{\sqrt{j!}}\,\frac{1}{\sigma}\,\varphi^{(j)}\left(\frac{\mu}{\sigma}\right),$$

where $\varphi^{(j)}$ denotes the j-th derivative of the standard Gaussian density with respect to its argument. It is well known that the convolution of two Gaussian random variables is also a Gaussian random variable whose variance is the sum of the original variances, and that convolution commutes with differentiation.

In terms of the change of variables $s = x/\sigma$, we see that

$$\begin{aligned}
&\left(\varphi * \left(\frac{1}{\sigma}\varphi\left(\frac{\cdot}{\sigma}\right) \cdot H_j\left(\frac{\cdot}{\sigma}\right)\right)\right)(x) \\
&\quad = \frac{1}{\sqrt{1+\sigma^2}} \frac{1}{\sqrt{j!}} \left(\frac{\partial}{\partial s}\right)^j \varphi\left(\frac{\sigma}{\sqrt{1+\sigma^2}}\, s\right) \\
&\quad = \frac{1}{\sqrt{1+\sigma^2}} \frac{1}{\sqrt{j!}} \left(\frac{\sigma}{\sqrt{1+\sigma^2}}\right)^j \varphi^{(j)}\left(\frac{x}{\sqrt{1+\sigma^2}}\right).
\end{aligned}$$

Plugging this expression into our previous formula, we find that

$$\begin{aligned}
&\int \nu_j(\mu_1)\, P_g(\mu_1, \mu_2)\, \nu_k(\mu_2)\, d\mu_1 d\mu_2 \\
&\quad = \frac{1}{\sqrt{1+\sigma^2}} \left(\frac{\sigma}{\sqrt{1+\sigma^2}}\right)^{j+k} \int \frac{\varphi^{(j)}\left(\frac{x}{\sqrt{1+\sigma^2}}\right) \varphi^{(k)}\left(\frac{x}{\sqrt{1+\sigma^2}}\right)}{\sqrt{j!}\sqrt{k!}\, \varphi\left(\frac{x}{\sqrt{1+\sigma^2}}\right)} dx \\
&\quad = \left(\frac{\sigma}{\sqrt{1+\sigma^2}}\right)^{j+k} \int H_j\left(\frac{x}{\sqrt{1+\sigma^2}}\right) H_k\left(\frac{x}{\sqrt{1+\sigma^2}}\right) \\
&\qquad\qquad \frac{1}{\sqrt{1+\sigma^2}} \varphi\left(\frac{x}{\sqrt{1+\sigma^2}}\right) dx \\
&\quad = \left(\frac{\sigma}{\sqrt{1+\sigma^2}}\right)^{j+k} \int H_j(x)\, H_k(x)\, \varphi(x)\, dx \\
&\quad = \left(\frac{\sigma}{\sqrt{1+\sigma^2}}\right)^{j+k} \delta(\{j=k\}),
\end{aligned}$$

because the Hermite polynomials as defined in (26) are orthonormal with respect to the standard Gaussian distribution. By Theorem 3, we thus conclude that the most favorable family is given by the first p Hermite polynomials. Moreover, again by Theorem 3, the relative efficiency coefficient of this family corresponds to the $p+1$-st eigenvalue of P_g, i.e., $\left(\sigma^2/\left(1+\sigma^2\right)\right)^j$. Because Γ_j was a most favorable family for density deconvolution, any other family of the form (2) will have worse relative efficiency.

Proof of Lemma 7

Continuing the argument our argument from Section 3.1, it remains to derive a lower bound for the minimax risk R^L among linear estimators in the Gaussian sequence model. We begin by noting that

$$\kappa^{J_C} \leq \mu_C < \kappa^{J_C+1}, \quad J_C = \left\lfloor \log(\mu_C) \,/\, \log(\kappa) \right\rfloor.$$

Thus, we can expand out (47) as

$$\begin{aligned}
C^2 &= \sum_{j=1}^{J_C} r_\sigma^{2j} \kappa^j \left(\mu_C - \kappa^j\right) \\
&= \mu_C \, \frac{\left(r_\sigma^2 \kappa\right)^{J_C+1} - r_\sigma^2 \kappa}{r_\sigma^2 \kappa - 1} - \frac{\left(r_\sigma^2 \kappa^2\right)^{J_C+1} - r_\sigma^2 \kappa^2}{r_\sigma^2 \kappa^2 - 1} \\
&= \left(r_\sigma^2 \kappa^2\right)^{J_C+1} \left(\frac{\mu_C \, \kappa^{-(J_C+1)}}{r_\sigma^2 \kappa - 1} - \frac{1}{r_\sigma^2 \kappa^2 - 1}\right) - \left(\frac{\mu_C \, r_\sigma^2 \kappa}{r_\sigma^2 \kappa - 1} - \frac{r_\sigma^2 \kappa^2}{r_\sigma^2 \kappa^2 - 1}\right) \\
&\leq B_{\sigma,\kappa}^2 \left(r_\sigma^2 \kappa^2\right)^{J_C+1} - 1, \quad B_{\sigma,\kappa}^2 = \frac{r_\sigma^2 \kappa \left(\kappa - 1\right)}{\left(r_\sigma^2 \kappa^2 - 1\right)\left(r^2 \kappa - 1\right)},
\end{aligned}$$

where the inequality on the third line hold whenever C (and thus also μ_C and J_C) are large enough. Thus, we conclude that

$$J_C + 1 \geq \frac{\log\left(\sqrt{C^2+1}/B_{\sigma,\kappa}\right)}{\log\left(r_\sigma \kappa\right)}.$$

We can also bound the risk R^L in (46) by

$$\begin{aligned}
R^L &= \sum_{j=1}^{J_C} r_\sigma^{2j} \left(1 - \frac{\kappa^j}{\mu_C}\right) \\
&\geq \sum_{j=1}^{J_C} r_\sigma^{2j} \left(1 - \frac{\kappa^j}{\kappa^{J_C+1}}\right) \\
&= \frac{r_\sigma^{2(J_C+1)} - r_\sigma^2}{r_\sigma^2 - 1} - \frac{\left(r_\sigma^2 \kappa\right)^{J_C+1} - r_\sigma^2 \kappa}{\kappa^{J_C+1}\left(r_\sigma^2 \kappa - 1\right)} \\
&\geq r_\sigma^{2(J_C+1)} \, \frac{r_\sigma^2 \left(\kappa - 1\right)}{\left(r_\sigma^2 - 1\right)\left(r_\sigma^2 \kappa - 1\right)} - \frac{r_\sigma^2}{r_\sigma^2 - 1}.
\end{aligned}$$

Plugging in our previous bound for J_C, we find that

$$\begin{aligned}
R_L &\geq \exp\left[\log\left(\frac{\sqrt{C^2+1}}{B_{\sigma,\kappa}}\right) \frac{2\log\left(r_\sigma\right)}{\log\left(r_\sigma\right) + \log\left(\kappa\right)}\right] \frac{r_\sigma^2 \left(\kappa - 1\right)}{\left(r_\sigma^2 - 1\right)\left(r_\sigma^2 \kappa - 1\right)} - \frac{r_\sigma^2}{r_\sigma^2 - 1} \\
&\geq \exp\left[\log\left(\frac{C}{B_{\sigma,\kappa}}\right) \frac{2\log\left(r_\sigma\right)}{\log\left(r_\sigma\right) + \log\left(\kappa\right)}\right] \frac{r_\sigma^2 \left(\kappa - 1\right)}{\left(r_\sigma^2 - 1\right)\left(r_\sigma^2 \kappa - 1\right)},
\end{aligned}$$

where the last inequality again holds for large enough C. Once paired with the 5/4 bound of Donoho et al. [1990] for Pinsker's constant, this bound yields the desired result.

Proof of Lemma 8

By the same argument as in the proof of 9, we can verify that

$$\begin{aligned}\lim_{M\to\infty}\lim_{n\to\infty}\mathbb{E}\left[L_n\left(\hat{g}^{(n,\,p)}\right)\right] &= \lim_{n\to\infty}\sum_{j=1}^{\infty}\mathbb{E}\left[Z_j-\gamma_j\right]^2\\ &= \sum_{j=1}^{p} r_\sigma^{2j}+\sum_{j=p+1}^{\infty}\gamma_j^2.\end{aligned}$$

The above expression is largest if the signal concentrates in the $p+1$-st Hermite coefficient, i.e., $\tau(\mu)=\gamma_{p+1}H_j(\mu/\sigma)$, yielding

$$\begin{aligned}\limsup_{n\to\infty}\left\{\sup_{\tau\in A^{\sigma}_{\kappa,\,C}} L_n\left(\hat{g}^{(n,\,p)}\right)\right\} &= \sum_{j=1}^{p} r_\sigma^{2j}+C^2\kappa^{-2(p+1)}\\ &\le \frac{1}{r_\sigma^2-1}r_\sigma^{2(p+1)}+C^2\kappa^{-2(p+1)}.\end{aligned}$$

Plugging in the choice for p specified in (35) and assuming that C is large enough that $\lceil\log(C/\sigma)/\log(r_\sigma\kappa)\rceil-1\ge 2$, we get that

$$\begin{aligned}\limsup_{n\to\infty}\left\{\sup_{\tau\in A^{\sigma}_{\kappa,\,C}} L_n\left(\hat{g}^{(n,\,p)}\right)\right\} &\le \sigma^2 r^{2\left(\frac{\log(C/\sigma)}{\log r_\sigma\kappa}+1\right)}+C^2\kappa^{-2\left(\frac{\log(C/\sigma)}{\log r_\sigma\kappa}\right)}\\ &= \left(1+r^2\right)\sigma^{\frac{2\log(k)}{\log(r_\sigma k)}}C^{\frac{2\log(r_\sigma)}{\log(r_\sigma k)}},\end{aligned}$$

which is what we set out to show.

Proof of Lemma 9

Our proof proceeds in several parts. We begin by establishing a version of our result for a simpler finite-dimensional problem; the general statement then shows that the finite-dimensional case can uniformly approximate our problem of interest.

A Finite-Dimensional Model For some $J\in\mathbb{N}$, suppose that τ is known to lie in an ellipse $A^{\sigma}_{\kappa,\,C}(J)$ defined by

$$\tau(\mu)=\sum_{j=1}^{J}\gamma_j H_j\left(\frac{\mu}{\sigma}\right),\quad \gamma\in\ell_2^{\kappa}(C,\,J):=\left\{\gamma':\sum_{j=1}^{J}\kappa^{2j}\gamma_j^2\le C^2\right\},$$

and that we only consider estimators over the set $\Lambda^{\sigma}_{\kappa}(J)$ defined by

$$\hat{g}^{(n)}(\mu)=g_\sigma^M(\mu)\exp\left[\frac{1}{\sqrt{n}}\sum_{j=1}^{J}\hat{\gamma}_j H_j\left(\frac{\mu}{\sigma}\right)-\psi_n(\hat{\gamma})\right],\quad \sum_{j=1}^{J}\kappa^{2j}\hat{\gamma}_j^2<L^2.$$

Our first task is to show that the the minimax risk over this finite-dimensional parametric class can, for large M, be well-approximated by minimax risk of the analogous finite Gaussian problem. We recall that, as usual, $g_\sigma^M(\mu)$ denotes the Gaussian density $g_\sigma(\mu)$ that has been "wrapped around" the interval $\Omega_M = [-M, M]$.

Convergence of the Likelihood Consider any parameter γ' whose induced tilting function satisfies $\tau_{\gamma'} \in A^\sigma_{\kappa, C}(J)$, and denote the resulting marginal density function by $f^{(n)}_{\gamma'} \propto \varphi * g_\sigma^M e^{\tau_{\gamma'}/\sqrt{n}}$. Because the basis functions H_j are all bounded (recall that we assume a compact domain $[-M, M]$), the log-likelihood $f^{(n)}_{\gamma'}(x)$ is uniformly Lipschitz in γ' for all x. Thus, we can use standard empirical process theory results to verify that the log-likelihood at γ', namely $\log(\prod_{i=1}^n f^{(n)}_{\gamma'}(X_i))$, is entirely determined by the score at the optimum γ [e.g., van der Vaart, 2000, Lemma 19.31]:

$$\sup_{\tau_{\gamma'} \in A^\sigma_{\kappa, C}(J)} \left| \sqrt{n} \log\left(\frac{\prod_{i=1}^n f^{(n)}_{\gamma'}(X_i)}{\prod_{i=1}^n f^{(n)}_{\gamma}(X_i)} \right) - (\gamma' - \gamma) \cdot \sum_{i=1}^n \nabla \log\left(f^{(n)}_{\gamma}(X_i) \right) \right| \to_p 0.$$

As γ is the true optimal parameter, we know that $\mathbb{E}_\gamma\left[\nabla \log\left(f^{(n)}_\gamma(X_i)\right)\right] = 0$; meanwhile, $n \operatorname{Var}_\gamma\left[\nabla \log\left(f^{(n)}_\gamma(X_i)\right)\right]$ converges to the Fisher information at $\gamma = 0$. Thus, by the central limit theorem, we can verify that

$$\sum_{i=1}^n \nabla \log\left(f^{(n)}_{\gamma}(X_i) \right) \Rightarrow \mathcal{N}(0, V_{M,J}), \quad V_{M,J} := \operatorname{Var}_0\left[\nabla \log\left(f_0(X_i)\right)\right];$$

Thus, we conclude that the log-likelihood in favor of γ' relative to γ is asymptotically equivalent in distribution to the log-likelihood arising from the Gaussian experiment where we observe $Z_M \sim \mathcal{N}(\gamma, V_{M,J})$ and want to recover $\gamma \in \ell_2^\kappa(C, J)$.

Convergence of the Loss Similarly, we can verify that the density estimation loss at γ' given the true parameter value γ satisfies

$$\begin{aligned} L_n\left(g^{(n)}_{\tau_{\gamma'}}\right) = n \int_{\Omega_M} & g_\sigma^M(\mu)\, e^{\frac{1}{\sqrt{n}}\gamma \cdot H_{(1:J)}\left(\frac{\mu}{\sigma}\right) - \psi_M\left(\frac{\gamma}{\sqrt{n}}\right)} \\ & \left(\frac{1}{\sqrt{n}}(\gamma - \gamma') \cdot H_{(1:J)}\left(\frac{\mu}{\sigma}\right) - \psi_M\left(\frac{\gamma}{\sqrt{n}}\right) + \psi_M\left(\frac{\gamma'}{\sqrt{n}}\right) \right) d\mu, \end{aligned}$$

and so

$$\begin{aligned} \lim_{n\to\infty} L_n\left(g^{(n)}_{\tau_{\gamma'}}\right) &= \int_{\Omega_M} \left((\gamma - \gamma') \cdot \left(H_{(1:J)}\left(\frac{\mu}{\sigma}\right) - \psi'_M(0) \right) \right)^2 g_\sigma^M(\mu)\, d\mu \\ &= (\gamma' - \gamma)^\top W_{M,J}(\gamma' - \gamma), \\ \text{with } W_{M,J} &:= \int_{\Omega_M} \left(H_{(1:J)}\left(\frac{\mu}{\sigma}\right) - \psi'_M(0) \right)^{\otimes 2} g_\sigma^M(\mu)\, d\mu, \end{aligned}$$

where $H_{(1:J)}(\cdot)$ is a vector obtained by stacking the first J Hermite functions. Moreover, this convergence is uniform in γ'. Thus, for large n, our re-scaled deviance loss is asymptotically equal to

$$L_{M,J}(\gamma') := (\gamma' - \gamma)^\top W_{M,J}(\gamma' - \gamma).$$

Convergence of the Minimax Risk Given the convergence results for the log-likelihood and for the loss established above, we should expect the statistical problem of density estimation to be asymptotically equivalent to finding the mean of a Gaussian vector with variance $V_{M,J}$ under loss $L_{M,J}$. To establish this formally, we can use least-favorable priors. Let $\hat{\gamma}^*_{M,J}$ be a minimax estimator for γ in the Gaussian model with variance $V_{M,J}$ under loss $L_{M,J}$, subject to $\gamma \in \ell_2^\kappa(C, J)$. By checking the conditions of Wald [1945], we can verify that $\hat{\gamma}^*_{M,J}$ is unique and that it is also Bayes for a least-favorable prior $\pi^*_{M,J}$; moreover, the risk of $\hat{\gamma}^*_{M,J}$ is constant over the support of $\pi^*_{M,J}$. Finally, because $L_{M,J}$ is quadratic, $\hat{\gamma}^*_{M,J}$ is the posterior mean for this least-favorable prior.

Now, let $\hat{\gamma}^{(n,\pi^*)}_{M,J}$ be the posterior mean for γ in the density estimation problem with n samples, where γ has a prior $\pi^*_{M,J}$. By our previous results, we find that

$$\mathcal{L}\left(L_n\left(g^{(n)}_{\tau_{\hat{\gamma}(n,\pi^*)_{M,J}}}\right)\right) \Rightarrow \mathcal{L}\left(L_{M,J}\left(\hat{\gamma}^*_{M,J}\right)\right),$$

and that the risk of this estimator is asymptotically constant over the support of $\pi^*_{M,J}$; thus, this estimator is asymptotically minimax over the support of $\pi^*_{M,J}$. Finally, the risk of this estimator is never asymptotically worse outside the support of $\pi^*_{M,J}$, and so $\hat{\gamma}^{(n,\pi^*)}_{M,J}$ is in fact asymptotically minimax for our finite-dimensional density estimation problem.

Taking Limits To move from the compact interval Ω_M to the real line we observe that, for any fixed J, we know that the Hermite functions are orthogonal over the whole real line:

$$\int_{\mathbb{R}} H_{(1:J)}\left(\frac{\mu}{\sigma}\right) g_\sigma(\mu)\ d\mu = 0,$$
$$\int_{\mathbb{R}} H_{(1:J)}\left(\frac{\mu}{\sigma}\right)^{\otimes 2} g_\sigma(\mu)\ d\mu = I_{J\times J},$$

and so $\lim_{M\to\infty} W_{M,J} = I_{J\times J}$. Moreover, by the same argument as used in the proof of Theorem 5, we find that $V_{M,J}$ converges to the leading $J \times J$ sub-matrix of the covariance matrix defined in (43). Meanwhile, for a fixed M, we can use the uniform approximability result from Corollary 4 to verify that the minimax risk of density estimation converges as $J \to \infty$; a similar argument holds in the Gaussian case. Because the minimax risks for both problems match for every finite J, they must thus also match as $J \to \infty$. Now, because we have established convergence both when J goes to infinity given a fixed M and when M goes to infinity given a fixed J, we conclude that the joint limit $M, J \to \infty$ is well-defined and does not depend on the order in which we take the limits; this implies the desired result.

Proof of Lemma 10

Asymptotic normality follows directly from the central limit theorem; we only need to verify the moments. Now, using similar arguments as in the proof of Corollary 4, we can verify that the limiting covariance of the $Z_j^{(n)}$ does not depend on τ, and that

$$\begin{aligned}
\lim_{n\to\infty} \operatorname{Cov}\left[Z_j^{(n)}, Z_{j'}^{(n)}\right] &= \rho_j^{-2} \int_\Omega U_j(x) U_{j'}(x) f_0(x)\, dx \\
&= \rho_j^{-2} \int_{\Omega^3} \frac{K(x, \mu_1)\, T_j(\mu_1)\, g_0(\mu_1)\; K(x, \mu_2)\, T_{j'}(\mu_2)\, g_0(\mu_2)}{f_0(x)}\, d\mu_1\, d\mu_2\, dx \\
&= \rho_j^{-2} \int_{\Omega^2} \zeta_{j+1}(\mu_1)\; P_{g_0}(\mu_1, \mu_2)\; \zeta_{j'+1}(\mu_2)\; d\mu_1\, d\mu_2 \\
&= \delta_{\{j=j'\}}\, \rho_j^{-1},
\end{aligned}$$

where P_{g_0} is the linear operator defined in (21) and the ζ_j are its eigenvectors as defined in Theorem 3; recall that $\rho_j = \lambda_{j+1}(P_{g_0})$. Meanwhile, the T_j are centered such that $\mathbb{E}[T_j(\mu)] = 0$, and so

$$\begin{aligned}
\frac{1}{\sqrt{n}} \mathbb{E}_0\left[Z_j^{(n)}\right] &= \int_\Omega U_j(x) f_0(x)\, dx \\
&= \int_\Omega K(x, \mu) \int_\Omega T_j(\mu)\, g_0(\mu)\; d\mu\, dx = 0.
\end{aligned}$$

Thus, we can verify that

$$\begin{aligned}
\lim_{n\to\infty} \mathbb{E}\left[Z_j^{(n)}\right] &= \lim_{n\to\infty} \sqrt{n}\, \rho_j^{-1} \int_\Omega U_j(x)\, f_{\tau/\sqrt{n}}(x)\, dx \\
&= \rho_j^{-1} \int_\Omega U_j(x) \left[\frac{\partial}{\partial \varepsilon} f_{\varepsilon\tau}(x)\right]_{\varepsilon=0} dx \\
&= \rho_j^{-1} \int_{\Omega^3} \frac{K(x, \mu_1)\, T_j(\mu_1)\, g_0(\mu_1)\; K(x, \mu_2) \left[\frac{\partial}{\partial \varepsilon} g_{\varepsilon\tau}(\mu_2)\right]_{\varepsilon=0}}{f_0(x)}\, d\mu_1\, d\mu_2\, dx \\
&= \rho_j^{-1} \int_{\Omega^2} T_j(\mu_1)\; \sqrt{g_0(\mu_1)}\, P_{g_0} \sqrt{g_0(\mu_2)}\, \tau(\mu_2)\; d\mu_1\, d\mu_2 \\
&= \int_\Omega T_j(\mu)\; \tau(\mu)\; g_0(\mu)\; d\mu,
\end{aligned}$$

where the last equality follow from the spectral theorem because $T_j\sqrt{g_0}$ is an eigenfunction of P_{g_0} with eigenvalue ρ_j.